

Universidade Federal do Rio Grande do Sul

Escola de Pós graduação em Engenharia Elétrica

Hardware image processing for high speed, high
resolution, large array detectors for scientific images.

Doctorate proposal

Dionisio Doering

Candidate

Prof. Dr. Carlos Eduardo Pereira

Adviser

Porto Alegre, December 2010.

Table of Contents

Introduction.....	3
Motivation.....	3
System overview	4
Sensors	5
Hardware architectures.....	8
Algorithms	8
State of the art	9
PhD project summary.....	11
Bibliography.....	13
Attachment A	14
Course Description	14
Course Schedule	14
Participating sites	15
Course certificate	16

Introduction

Image sensors have been in use for many years on consumer and scientific applications. The evolution of such sensors made them larger, faster and cheaper. This has allowed people to use them on a daily basis, for example, on surveillance systems, hand held camcorders, cell phones, etc. This evolution makes also the scientific detectors larger in terms of number of pixels (i.e. more than 1M pixels) and high speed (i.e. more the 30 frames per second (fps)). One consequence of these large and high speed detectors is the generation of large data volumes. This large data set has to be sent somewhere, for example, to a display or storage unit. In the latter this means that more media is necessary to store this data. Transmission buses and storage devices are limited resources and solutions need to be developed to improve their capabilities to cope with this data volume.

In the consumer world the large data volume problem is being solved by applying compression algorithms (with data loss) implemented directly in the sensor. This solution keeps the bandwidth of the detector smaller. For applications where information cannot be lost (i.e. scientific or medical) the same solution cannot be applied, and new ones need to be created. Part of the solution can be done through real time image processing implemented in the sensor chip, or in the surrounding system, or both. There are many aspects that need to be optimized to create a system for this application. For example, there are several architectures that can be used which include Field Gate Programmable Arrays (FPGA), Digital Signal Processors (DSP), micro controllers (μ C), Central Processing Units (CPU), Graphical Processing Units (GPU) or a mix of them creating a distributed system. Another important aspect for this system is the algorithms that will be implemented in each node of the system architecture. For the nodes where CPU are present then it is also necessary to look into operational systems and which ones are more suitable.

The present project proposes to make use of the existing model driven methodologies for real time distributed systems (developed at the GCAR (Grupo de Controle Automacao e Robotica), under Prof. Carlos Eduardo Pereira's research topics) to study architectures, evaluate algorithms, possible target hardware platforms and propose a frame work for the development of efficient real time image processing systems for high speed, high resolution large detector arrays for scientific application.

Motivation

There are many innovations of detector development nowadays and the Lawrence Berkeley National Lab through the ICES group (Integrated Circuit and Electronics Systems, from Engineering division) is one of them[6,9,10]. This group develops new detectors in several technologies (for example CCD, CMOS and SOI) for scientific applications on X-ray image detector and electron microscopy. Among the developments of new detectors there are a 1Mpixel frame store CCD running at 200 fps; a 1Mpixels and a 4Mpixel CMOS sensors running at 400 fps. Bigger and faster cameras for use in scientific applications are in the road map of this group but even the existing cameras are already a real example

of the challenges that high speed, high resolution large detectors create. The existing and future detectors developed at LBNL are part of the motivation to study the new challenges that high speed real time hardware image processing are imposing on their readout systems.

The LBNL detectors will also serve as a test bed for the proposed framework that will be developed during this research. Currently the most used approach to deal with data from the image detectors is to send the data to a storage system for offline processing. This approach is not practical for the new cameras and it is not efficient in terms of use of backplane buses and storage units.

The research developed at the GCAR (Grupo de Controle Automacao e Robotica) under Prof. Dr. Carlos Eduardo Pereira in real time distributed systems provide a set of supporting knowledge and expertise that can be applied on the development of a framework that can generate efficient systems to be used with the detectors develop at LBNL.

The study of hardware image processing for high speed cameras, high resolution and large detector array with the goal of developing better solutions is a very challenging research topic on both theoretical and practical aspects. This topic is not limited to the applications being studied at LBNL though; there are many other fields that may benefit from this research such as autonomous robot guidance systems, surveillance systems, visual quality inspections for manufactures, etc. As these new detectors become more available then more applications will benefit from the outcomes of this research as well.

System overview

A high speed, high resolution large detector system has similar building blocks than many other imaging systems. The challenges that the former system has and the latter doesn't are caused by the combination of high speed, high resolution large detector array and the use on compression algorithms that don't lose information concurrently. Next figure shows a block diagram of an image system. In the following sections there is a description of each one of these building blocks.

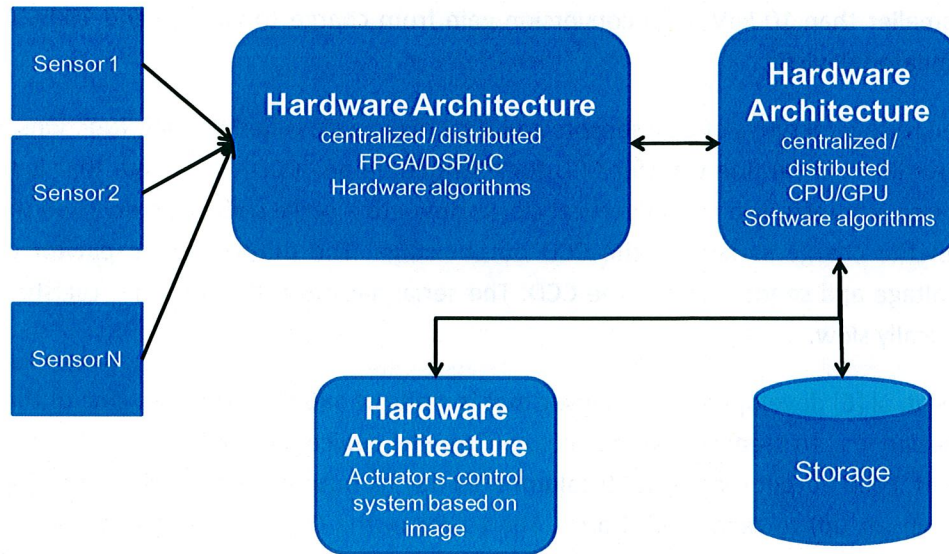


Figure 1 - Building block of an image system.

Sensors

Charge Coupled Device (CCD) was invented by Boyle and Smith [7] in 1970. Since then this devices have been studied and improved in terms of its physics, fabrication and operation. The CCD sensor has three main functions, which are (1) charge collection, (2) charge transfer and (3) charge conversion into a measurable voltage. The CCD array requires a series of clocks and bias voltage to control it. These signals are provided by the system that surrounds the sensor. Figure 2 shows a schematic view of a CCD device (a) and its serial shift register (b).

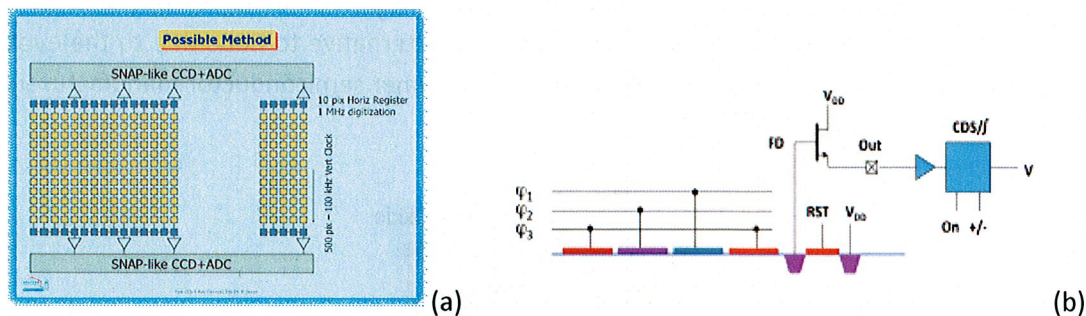


Figure 2 – Schematic view of (a) the LBNL FCCD and (b) serial shift register of a CCD device.

CCDs have been extensively used for X-rays detection. The detection can be done in two modes, indirect detection (a fiber optic taper with a phosphor screen is used) or direct detection (the X-rays interact direct with the silicon that forms the sensor). CCDs are a reasonable choice when one wants to image X-rays using direct detection because this device has very low noise, good efficiency for X-rays

with energy smaller than 10 keV, high conversion gain from charge to voltage and 100% fill factor for backside illuminated devices.

There are several types of CCD sensors and they may be classified by their functions, such as, full frame CCD, interline CCD, frame transfer CCD for example. In any CCD the charge that is generated in the sensitive area is transferred by the vertical clocks down to a serial shift register. Then the horizontal clocks transfer the charge serially to the CCD output stage. The output gate capacitor converts the charge into voltage and sends it out of the CCD. The serial nature of this last shift register makes CCD devices intrinsically slow.

Denes et al. [6] developed a CCD sensor that is made of 480 by 480 pixels. One of the techniques used in this design was to create an almost column parallel device. In it, 48 outputs in each side of the CCD are present. Each output reads out 10 columns. By having outputs on both sides only 240 rows have to be shift down (or up) to read a full frame. The traces in the CCD are made of poly-silicon material which has a resistivity that is about 1000 times bigger than the aluminum resistivity for example. Since CCDs are a huge capacitance array the maximum frequency that one can clock the CCD has a limit that is set by the R of the poly-silicon and by the C of the CCD array. Another technique used in the sensor described by Denes et al. is called metal strapping. In it a trace made of aluminum is implanted in parallel with the poly-silicon making the CCD charge transfer to be faster than most CCD available up until now. These new techniques allowed the Lawrence Berkeley National Laboratory (LBNL) FCCD (fast CCD) to run at about 200 fps.

New devices are being designed at LBNL. Among them a 1Mpixel frame store device (this device can also be used as a 2Mpixel full frame CCD). This device also runs at 200 fps. The data bandwidth for such device is 400MB/s. A system that is capable to handle this bandwidth is being designed for such camera. But even though it is possible to save raw data to a storage system it makes sense to perform some real time image processing to reduce the amount of data that is stored.

Complementary metal oxide semiconductor (CMOS) was patented in 1967 by Frank Wanlass (US patent 3,356,858) [7]. In 1990, CMOS emerged as a serious alternative to CCDs due to the evolution of the CMOS processing technology that was driven by the consumer semiconductor manufacturers. Figure 3 shows the schematic of a generic CMOS device.

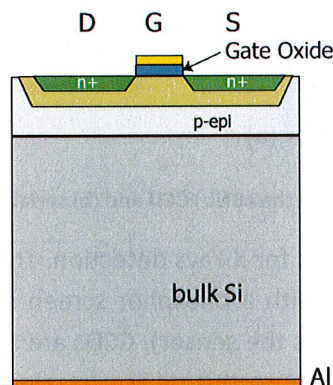


Figure 3 - Schematic of a generic CMOS device containing a single transistor with drain (D), source (S) and gate (G) electrodes. Electrodes are fabricated on a bulk silicon wafer, with the thin p-epitaxial layer serving as the site of electron-hole pair generation. Electrons are incident from above.

The processing evolution allowed the creation of commercial cameras but it was not until 2009 when radiation hard devices were developed [9-10] that CMOS cameras started to make part of direct detected electron microscopy. These new sensors besides the improvements on the radiation tolerance were also made to be readout at 400 fps. The first camera produced was a 1Mpixel camera called TEAM 1k. In this case the data bandwidth produced by this camera is 800MB/s.

The system to drive and digitize the data from the TEAM 1k sensor was done. The main challenge that is still under investigation is a reliable way to send this data into a storage system. One way to reduce this challenge is to introduce hardware image processing into the system and reduce the amount of data that needs to be stored by the storage system.

Another sensor that is being developed in CMOS technology for electron microscopy at LBNL is a 4Mpixel camera that also runs at 400 fps called TEAM 2k. In this case the camera data bandwidth is 3.2GB/s. In this case the challenge of sending data to storage and storage limitation is even more evident and hardware image processing is a very powerful tool to be used immediately.

Silicon-On-Insulator (SOI) refers to a process that is done using a layered wafer made of silicon-insulator-silicon. For image sensor the bottom layer is used as detector. The top silicon layer is used to implement the CMOS electronics. With this approach it is expected that cameras produced in this technology will achieve the noise and resolution levels of the CCD and the dense electronics of the CMOS cameras. Sensors using this technology are under investigation at the LBNL. As the SOI devices have similar readout electronics as the CMOS readout one can expect cameras with data bandwidth on the range of 400MB/s to 3.2GB/s. Figure 4 shows a generic schematic of an SOI device.

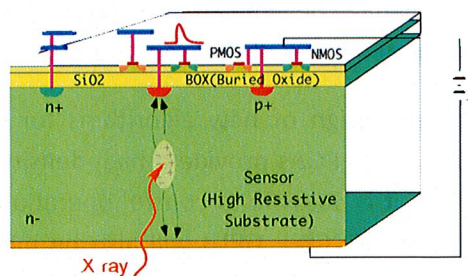


Figure 4 - Schematic of SOI device. Electron-hole pairs generated by x-ray or electron irradiation of the bulk material are collected at structures fabricated above and below the buried insulating layer.

Optical images are widely used for gathering scientific images. These camera sensors can generate color or monochromatic images. In the case of color images three physical pixels are needed to generate one image element. In front of each pixel a filter for red or green or blue is added allowing for the generation of RGB images. Both types of cameras can be used for optical microscopy, indirect detection for electron or X-ray images, machine vision for autonomous systems, object inspection, stereo image generation for 3D images to list just a few applications. Sensors for these applications can be done in all three technologies described earlier.

Hardware architectures

Hardware architectures for real time image processing can have many different arrangements. Below is a list of elements that can exist in them:

- Single/multiple sensors
- Centralized/distributed processing
- Hardware (FPGAs)/GPUs/PCs/super computers
- Reconfigurable hardware
 - Multipliers
 - Internal memory
 - Block/distributed RAM
 - Embedded systems

The camera systems developed at LBNL are based on Field Programmable Gate Array (FPGA) which is a configurable hardware that was invented by Ross Freeman and Bernard Vonderschmitt in 1985. These programmable chips consist of a large set of configurable logic blocks, interconnections, memory and specific hardware for some specific functions such as digital clock manager, high speed transceiver and DSP blocks. This is one example of system architecture and it was chosen to handle the present sensors.

These devices are programmed by hardware description language HDL. One example of such languages is VHDL which is an acronym for VHSIC that means very high speed integrated circuit, hardware description language [8]. This language permits the modeling of digital systems in several levels ranging from simple gate level logic to complete systems. The first version of VHDL 7.2 was released to the public in 1985. The development of this language was financed by the Department of Defense of the United States. This language became an IEEE standard in 1987 and it was revised in 1993, 2000, 2002, 2006 and 2008. The current version of the standard is called IEEE 1026-2008.

The proposed project will cover design of new algorithms for high speed scientific images for hardware implementation because [5] i) FPGAs provide a high-density for arithmetic logic at relatively short design cycles and (ii) in FPGAs, it is possible to control operations at bit level to build specialized data paths. Other architectures (DSP, uC, CPU, GPU) will also be modeled and the use of them in a distributed environment can then be simulated for evaluation of efficient use.

The hardware architecture also can contain a set of actuators. In systems where this is true then a closed loop control algorithm can be implemented. According to xxx [] a data acquisition system is a special case of a control system when there are no actuators. At this point, the actuator building block as well as any control algorithms is beyond the scope of this proposal.

Algorithms

The development of an efficient system through a frame work for efficient hardware image processing system needs also to cover the algorithms that will be used in this system. This is important because each algorithm may have a more suitable place where it should be executing leading to an efficient system. Below are a few examples of algorithms that are important for this application:

- Transforms
 - DCT
 - DWT
 - Wavelets
- Data compression
- Single/multiple frame analyses
- Image fusion
- Panning; 3D; Stereo; super resolution
- Fixed/moving sensors
- Single photon (electron/X-rays) information extractions

State of the art

The papers cited below represent some of the aspects that justify the importance of the proposed subject for this PhD thesis. This list is by far no complete but it is already a very strong source to justify this theses.

To be useful and effective, an engineering model must possess, to a sufficient degree, the following five key characteristics. The most important is abstraction. A model is always a reduced rendering of the system that it represents. By removing or hiding detail that is irrelevant for a given viewpoint, it lets us understand the essence more easily. Considering the steady demand for ever-more sophisticated functionality from our software systems, abstraction is almost the only available means of coping with the resulting complexity. The second key characteristic is understandability. It isn't sufficient just to abstract away detail; we must also present what remains in a form (for example, a notation) that most directly appeals to our intuition. Understandability is a direct function of the expressiveness of the modeling form used (expressiveness is the capacity to convey a complex idea with little direct information). A good model provides a shortcut by reducing the amount of intellectual effort required for understanding. One reason why programs are not particularly expressive, even when based on languages that support sophisticated abstractions, is that they require too much detailed parsing of text to be properly understood. Classical programming statements assault the reader with a profusion of syntactical detail assembled according to intricate lexical rules. The amount of information that must be absorbed and recognized to understand linear programs is enormous and requires significant intellectual effort. The third key characteristic of useful models is accuracy. A model must provide a true-to-life representation of the modeled system's features of interest. Fourth is productiveness'. You should be able to use a model to correctly predict the modeled system's interesting but non-obvious

properties, either through experimentation (such as by executing a model on a computer) or through some type of formal analysis. Clearly, this depends greatly on the model's accuracy and modeling form. For instance, a mathematical model of a bridge is much better at predicting the maximum allowable load on a bridge than even a very precise and detailed scale model constructed out of balsa wood. Finally, a model must be inexpensive—that is, it must be significantly cheaper to construct and analyze than the modeled system. [14]

Kalomiros and Lygouras [1] express the need to process images faster in order to be able to use them in real time. In this work they suggest an implementation of a auxiliary board that receives images from a pc, process them in hardware based on FPGA and send the processed data back to the pc for its use. One of the reasons for this approach justified by the authors is the difficulty of connecting cameras directly to the processing board. As a draw back their system spends a substantial amount of time on image transfer via USB port to the processing card.

Perez-Vidal and Garcia [2] present an implementation of computation expensive prediction filters based on image processing. In this work they show that an implementation based on FPGA parallel algorithm was 8 to 14 times faster than the pc based solution. Their experimental example didn't warrant the use of FPGA system instead of PC-based because both of them were faster than their system requirement. But they did justify that this approach is the best in cases where (1) There is more than one visual feature: usually several hundreds of features are used in visual servoing tasks. (2) The prediction filter has been computed for two dimensions but for position based visual servoing this computation must be extended to 6 dof. (3) Nowadays, high speed cameras are becoming more extended and reduce image acquisition ratio from 40 ms (25 fps) to 4 ms (250 fps). (4) More than one camera can be used (e.g. stereo vision). Difficulties found in the implementation are: mathematical operations must be programmed taking into account the overflow and round-off of signed fixed-point variables; matrices and vectors must be created and loaded; several clocks must be synchronized; Xilinx VHDL macros must be utilized to communicate with PC; and access to DDR2 SRAM must be programmed directly in VHDL using Memory Interface Generator (MIG).

Massoudina et al. [3] describe an FPGA based implementation of a compression algorithm. This algorithm uses wavelet transform. The compression is done by applying DCT to the base image (the output of the low pass filter). The authors also present detailed information on how to implement the masks in hardware using FIFOS and custom made multipliers base on logic instead of real multipliers. This is necessary to cope with the limited resources available in the FPGA parts. The results demonstrate that this approach is capable to handle 16Mpixel images with rate of up to 35 frames per second.

Karmani at al. [4] propose a new hardware architecture for 2D scan-based wavelet watermarking. The hardware is aimed for broadcast and HDTV. In their system they make use of the FPGA resources to implement several instances in parallel for the filters used. The video protection is achieved by insertion of watermarks bank within the middle frequency of wavelet coefficients related to video frames. This paper focus on the selection of appropriate real time watermarking scheme and performing a trade off between the algorithmic aspects and the hardware implementation techniques.

The authors use a 256 by 256 images that will be computed as claimed by the authors with a severe time constraints fixed by the frequency of the video speed scanning of 25 images per second.

Chrysafis and Ortega [5] address the problem of low memory wavelet image compression. The authors focus their attention to the issue of whether both the wavelet transforms and the subsequent coding can be implemented in low memory without significant loss in performance. The authors use a line base approach where only a minimum amount of lines are kept in memory. The authors also proposed a novel context-based encoder which requires no global information and stores only a local set of wavelets coefficients.

Applications of image processing using FPGAs are becoming more popular. Applications such as image merging using alpha channel [12] or automated ceramic tile inspections [13] are example of such applications. The main reason for doing these tasks in hardware (meaning FPGA) and not software is the speed of processing that these systems are capable of handling (video rate for images of 2k by 2k pixels) without overloading the computer system.

PhD project summary

The present project proposes to make use of the existing model driven methodologies for real time distributed system developed at the GCAR (Grupo de Controle Automacao e Robotica), under Prof. Carlos Eduard Pereira's research topics, to study architectures, evaluate algorithms, possible target hardware platforms and propose a framework for the development of real time image processing applications.

Follow it is shown the work break down structure for this project and its estimated schedule.

1. Classes

Four classes were taken during the year of 2010 adding to 12 credits. It is assumed that all 24 credits from the master degree program taken by the candidate will be validated for the PhD program and this will then conclude all 36 credits required by the program. Following are the description of the classes

ELE00002 – Sistemas de automacao

ELE00012 – Barramentos Industriais

ELE00100 – Programacao de sistemas distribuidos

Many Core Processors - Summer School 2010 at the Virtual School of Computational Science and Engineering (Please see Attachment A for more details on this class).

2. Bibliographic study

3. Theoretical study of image processing architectures, algorithms, platforms and models.

4. Proposal of models, architecture and algorithms for real time image processing in the framework designed during this research project.

5. Implementation and validation of the framework.

6. Write PhD theses

Item description	2011				2012				2013				2014			
	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
Classes (completed in 2010)																
Bibliographic revision	X	X	X	X												
Theoretical study		X	X	X	X	X	X	X								
Models, architecture and algorithm proposal				X	X	X	X	X	X	X	X	X				
Implementation and validation of the proposed framework for real time image processing						X	X	X	X	X	X	X	X	X		
Write theses															X	X

Bibliography

- [1] J.A. Kalomiros, J. Lygouras, Design and evaluation of a hardware/software FPGA-based system for fast image processing, *Microprocessors and Microsystems* 32 (2008) 95–106
- [2] C. Perez-Vidal, L. Gracia, High speed filtering using reconfigurable hardware, *J. Parallel Distrib. Comput.* (2009), doi:10.1016/j.jpdc.2009.06.004
- [3] A. Masoudnia et al., Design and performance of a pixel-level pipelined-parallel architecture for high speed wavelet-based image compression, *Computers and Electrical Engineering* 31 (2005) 572–588, doi:10.1016/j.compeleceng.2005.07.005
- [4] S. Karmani et al., Efficient Hardware architecture of 2D-scan-based wavelet watermarking for image and video, *Computer Standards & interfaces*, 31 (2009) 801-811, doi:10.1016/j.csi.2008.09.015
- [5] C. Chrysafis and A. Ortega, Line-Based, Reduced Memory, Wavelet Image Compression. *IEEE transactions on Image processing*, vol.9, No 3, 2000.
- [6] P. Denes, D. Doering, H. A. Padmore, J.-P. Walder, and J. Weizeorick, A fast, direct x-ray detection charge-coupled device, *REVIEW OF SCIENTIFIC INSTRUMENTS* 80, 1, 2009, #DOI: 10.1063/1.3187222
- [7] Holst, C. G., Lomheim, T. S., *CMOS/CCD Sensors and Camera Systems*. JCD Publishing. ISBN 9780970774934. Chapters 1 and 3.
- [8] Bhasker, J. *VHDL Primer*, 3rd edition, ISBN 0-13096575-8
- [9] Marco Battaglia, Devis Contarato, Peter Denes, Dionisio Doering, Velimir Radmilovic, CMOS pixel sensor response to low energy electrons in transmission electron microscopy, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Volume 605, Issue 3, 1 July 2009, Pages 350-352
- [10] Marco Battaglia, Devis Contarato, Peter Denes, Dionisio Doering, Piero Giubilato, Tae Sung Kim, Serena Mattiazzo, Velimir Radmilovic, Sarah Zalusky, A rad-hard CMOS active pixel sensor for electron microscopy, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Volume 598, Issue 2, 11 January 2009, Pages 642-649
- [11] Anderson, J.; Brito, R.; Doering, D.; Hayden, T.; Holmes, B.; Joseph, J.; Yaver, H.; Zimmermann, S.; Data Acquisition and Trigger System of the Gamma Ray Energy Tracking In-Beam Nuclear Array (GRETINA). *Nuclear Science, IEEE Transactions on*, Volume 56, Issue 1, Part 2, Feb. 2009 Page(s):258 - 265 DOI 10.1109/TNS.2008.2009444.
- [12] Xinxi Zhang; Yong Li; Jinyang Wang; Yulin Chen; Design of high-speed image processing system based on FPGA *Electronic Measurement & Instruments*, 2009. ICEMI '09. 9th International Conference on, 16-19 Aug. 2009 Page(s):4-65 - 4-69, Digital Object Identifier 10.1109/ICEMI.2009.5274670
- [13] Hocenski, Z.; Aleksi, I.; Mijakovic, R.; Ceramic tiles failure detection based on FPGA image processing, *Industrial Electronics*, 2009. ISIE 2009. IEEE International Symposium on 5-8 July 2009 Page(s):2169 – 2174 Digital Object Identifier 10.1109/ISIE.2009.5219911
- [14] SELIC, B. *The Pragmatics of Model-Driven Development*. IEEE Software, Los Alamitos, v.20, n.5, p.19–25, 2003a.

Attachment A

Course Description

The Proven Algorithmic Techniques for Many-core Processors has the objective to teach the most commonly used algorithm techniques and computational thinking skills needed for many-core programming, especially the simple ones! Specifically, to understand many-core hardware limitations and constraints; desirable and undesirable computation patterns; commonly used algorithm techniques to convert undesirable computation patterns into desirable ones.

For more information please visit <http://groups.google.com/group/vscse-many-core-processors-2010>

Course Schedule

VSCSE Many-core Processors 2010 (Pacific Daylight Time)						
Aug 1 - 7 2010						
Sun 8/1	Mon 8/2	Tue 8/3	Wed 8/4	Thu 8/5	Fri 8/6	Sat 8/7
	Proven Algorithmic Techniques for Many-core Processors with local activities in the western half of the US					
6am						
7am						
8am	8 - 9:30 Computational Thinking for Many-core Computing	8 - 9:30 Cut-off and Binning for Regular Data Sets	8 - 9:30 Keynote 2: Fermi and the Future of GPU Computing	8 - 9:30 Keynote 3: Multiplying speedups: fast algorithms on GPUs	8 - 8:45 Dealing with Dynamic Data	
9am	9:30 - 10:30 Break/Lab 0	9:30 - 10:30 Break/Lab 1	9:30 - 10:30 Break/Lab 2	9:30 - 10:30 Break/Lab 3	8:45 - 9:30 Directions for Further	
10am					9:30 - Closing Remarks	
11am	10:30 - 12p Scatter-to-Gather Transformation for Scalability	10:30 - 12p Data Layout for Grid Applications	10:30 - 11:45 PDE Solver Techniques	10:30 - 11:15 HPC Applications	11:15 - 12p Dealing with Dynamic Data	
12pm	12p - 1p Lunch	12p - 1p Lunch	12p - 1p Lunch	12p - 1p Lunch		
1pm	1p - 2:30p Loop Blocking and Register Tiling for Locality	1p - 2:30p Keynote 1: Algorithm Design for GPU Computing	1p - 2:30p Dealing with Non-uniform Data Distribution	1p - 2p Lab Discussion		
2pm	2:30p - 3:30p Lab 1	2:30p - 3:30p Lab 2	2:30p - 3:30p Lab 3	2p - 3:30p Lab 3		
3pm						
4pm						

Course Presentations - VSCSE Many-core Processors 2010 | Google Groups

http://groups.google.com/group/vscse-many-core-processors-2010/web/course-presentations

Google Groups will no longer be supporting the Pages and Files features. Starting January 13, you won't be able to upload new content, but you will still be able to view and download existing content. See [this announcement](#) for more information and other options for storing your content.

Course Presentations

See the [course schedule](#) for the precise times of each presentation. Speaker [biographies](#) are posted on a separate page.

Monday, August 2

- Lecture 1: Introduction to Computational Thinking for Many-core Computing (Wen-mei Hwu, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 2: Scatter-to-Gather Transformation for Scalability (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 3: Loop Blocking and Register Tiling for Locality (Wen-mei Hwu, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Tuesday, August 3

- Lecture 4: Cut-off and Binning for Regular Data Sets (Wen-mei Hwu and John Stone, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 5: Data Layout for Grid Applications (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Keynote 1: Algorithm Design for GPU Computing (Michael Garland, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Wednesday, August 4

- Keynote 2: Fermi and the Future of GPU Computing Technology (David Kirk, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 6: PDE Solver Techniques (Jonathan Cohen, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 7: Dealing with Non-uniform Data Distribution (Wen-mei Hwu, University of Illinois; David Kirk, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Thursday, August 5

- Keynote 3: Multiplying speedups: fast algorithms on GPUs (Lorena Barba, Boston University) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Guest Lecture: Accelerating HPC Applications with GPUs—Two Case Studies (Jeremy Meredith, Oak Ridge National Laboratory) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 8: Dealing with Dynamic Data Sets (Wen-mei Hwu, University of Illinois; David Kirk, NVIDIA) [\[video-part-1\]](#) [\[video-part-2\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Hands-on Lab Wrap-up Discussion (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Friday, August 6

- Lecture 9: Directions for Further Studies and Closing Remarks (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- gpuscomputing.net (Andrew Schuh, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Home

Discussions

Pages

Files

[About this group](#)

[Apply for group membership](#)

Sponsored links

GPU Computing Cluster
HPC Cluster with Intel® Xeon® CPUs and Tesla GPUs. Request a Quote!
[www.Appro.com](#)

University of Phoenix®
Online and Campus Degree Programs.
Official Site - Classes Start Soon.
[Phoenix.edu](#)

Group Info

Members: 158
Language: English
Group categories: Not categorized
[More group info >](#)

Recent pages and files

Course Presentation	Aug 24
Participating Sites	Aug 10
Hands-on Labs	Aug 9
Teaching Assistant	Aug 9
Solutions	Aug 6

Participating sites

Gmail Calendar Documents Reader Web more ▾

Google groups

VSCSE Many-core Processors 2010

Google Groups will no longer be supporting the Pages and Files features. Starting January 13, you won't be able to upload new content, but you will still be able to view and download existing content. See [this announcement](#) for more information and other options for storing your content.

Participating Sites

The following sites are fully participating in the *Proven Algorithmic Techniques for Many-core Processors* course:

- Center for Computation & Technology, Louisiana State University, Baton Rouge [\[local info\]](#) [\[photos\]](#)
- Institute for Data and High Performance Computing, Georgia Institute of Technology, Atlanta [\[local info\]](#) [\[photos\]](#)
- Institute for Digital Research and Education, University of California, Los Angeles [\[local info\]](#) [\[photos\]](#)
- National Center for Supercomputing Applications, Urbana, Illinois [\[local info\]](#) [\[photos\]](#)
- Northwestern University, Evanston, Illinois [\[local info\]](#) [\[photos\]](#)
- Ohio Supercomputer Center, Ohio State University, Columbus [\[local info\]](#) [\[photos\]](#)
- RENCI, Chapel Hill, North Carolina [\[local info\]](#) [\[photos\]](#)
- University of Iowa, Iowa City [\[local info\]](#) [\[photos\]](#)
- University of Michigan, Ann Arbor [\[local info\]](#) [\[photos\]](#)

The following sites will host remote presenters (with no audience):

- University of Minnesota Supercomputing Institute (Minneapolis)
- NVIDIA Corporation (Santa Clara, CA)

The following site is a receiving site only:

- Lawrence Berkeley National Laboratory [\[local info\]](#)

These sites are linked by a [high-definition videoconferencing technology](#), which provides a collaborative learning environment spanning geographical boundaries.

Version: [Current by Tom Scavo - Aug 10](#)

Home

Discussions

Pages

Files

[About this group](#)

[Apply for group membership](#)

Sponsored links

GPU Computing Cluster
HPC Cluster with Intel® Xeon® CPUs and Tesla GPUs. Request a Quote!
[www.Appro.com](#)

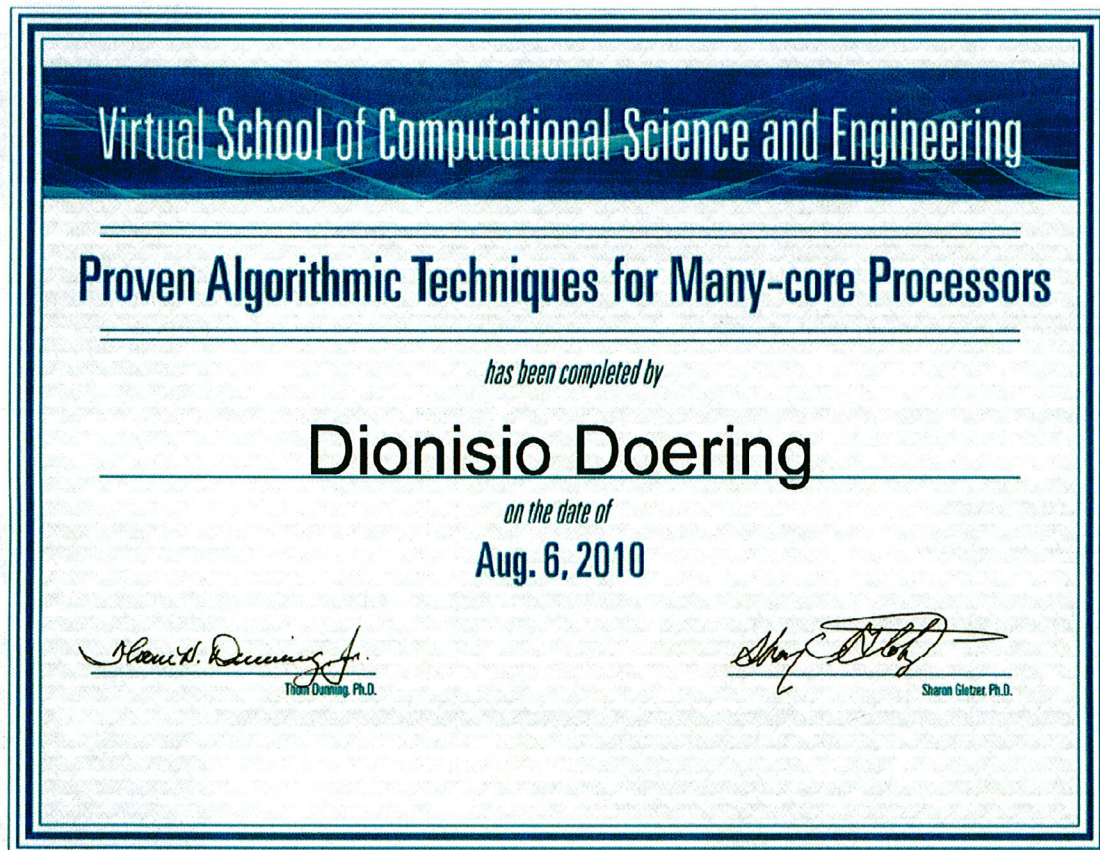
IVR made damn easy
Voxeo XML IVR platforms & hosting:
Free trials, extremely good support
[www.voxeo.com](#)

San Francisco Coupons
1 ridiculously huge coupon a day.
Get 50-90% off the Bay's best!
[www.Groupon.com/San-Francisco](#)

Group Info

Members: 158

Course certificate





Serviço Público Federal
UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
ESCOLA DE ENGENHARIA
DEPARTAMENTO DE ENGENHARIA ELÉTRICA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

TERMO DE COMPROMISSO

Através do presente Termo de Compromisso, firma-se o comprometimento de orientação do candidato Dionisio Doering, aprovado em processo seletivo para ingresso no Curso de Doutorado em Engenharia Elétrica da UFRGS.

O trabalho de doutorado do candidato será direcionado para a área de _____ e versará sobre Hardware image processing for high speed scientific images.

Porto Alegre, ____ de _____ de 200__.

Assinatura do Orientador



Serviço Público Federal
UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
ESCOLA DE ENGENHARIA
DEPARTAMENTO DE ENGENHARIA ELÉTRICA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

À Comissão de Pós-Graduação em
Engenharia Elétrica

Através do presente, solicito o reaproveitamento dos créditos obtidos em disciplina(s) em nível de pós-graduação, abaixo relacionada(s), conforme comprova o histórico escolar e a(s) respectiva(s) ementa(s) constante(s) da documentação que segue anexa:

Nome da disciplina	Período	Nº créditos	Conceito
1. Instrumentação (ELE00001)	03/01	03	A
2. Métodos Matemáticos para EE (ELE00017)	03/01	03	B
3. Introdução a Proc. Estoc. (ELE00019)	03/01	03	A
4. Proc. Avanç. Sinais (ELE00004)	03/02	03	A
5. Métodos Elem. Finitos (ELE00014)	03/02	03	A
6. Processamento de Imagens (ELE00005)	03/03	03	A
7. Trabalho Individual (ELE00040)	03/03	03	A
8. Tópicos Esp. Proc. Sinais (ELE00081)	03/03	03	A
9. Sistemas de automacao (ELE00002)	10/01	03	
10. Barramentos Industriais (ELE00012)	10/02	03	
11. Programacao de sis. Dist. (ELE00100)	10/03	03	
12. Many Core Processors	10/02	03	*

* Por favor ver detalhes da disciplina em anexo.

Porto Alegre, 10 de Dezembro de 2010.

Nome: DIONISIO DOERING

Parecer da Comissão de Pós-Graduação:

Ata da Comissão de Pós-Graduação nº _____, de ____/____/____.

Attachment A

Course Description

The Proven Algorithmic Techniques for Many-core Processors has the objective to teach the most commonly used algorithm techniques and computational thinking skills needed for many-core programming, especially the simple ones! Specifically, to understand many-core hardware limitations and constraints; desirable and undesirable computation patterns; commonly used algorithm techniques to convert undesirable computation patterns into desirable ones.

For more information please visit <http://groups.google.com/group/vscse-many-core-processors-2010>

Course Schedule

VSCSE Many-core Processors 2010 (Pacific Daylight Time)						
Aug 1 - 7 2010						
Sun 8/1	Mon 8/2	Tue 8/3	Wed 8/4	Thu 8/5	Fri 8/6	Sat 8/7
	Proven Algorithmic Techniques for Many-core Processors with local activities in the western half of the US					
6am						
7am						
8am	8 - 9:30 Computational Thinking for Many-core Computing	8 - 9:30 Cut-off and Binning for Regular Data Sets	8 - 9:30 Keynote 2: Fermi and the Future of GPU Computing	8 - 9:30 Keynote 3: Multiplying speedups: fast algorithms on GPUs	8 - 8:45 Dealing with Dynamic Data	
9am	9:30 - 10:30 Break/Lab 0	9:30 - 10:30 Break/Lab 1	9:30 - 10:30 Break/Lab 2	9:30 - 10:30 Break/Lab 3	8:45 - 9:30 Directions for Further	
10am	10:30 - 12p Scatter-to-Gather Transformation for Scalability	10:30 - 12p Data Layout for Grid Applications	10:30 - 11:45 PDE Solver Techniques	10:30 - 11:15 HPC Applications	9:30 - Closing Remarks	
11am	12p - 1p Lunch	12p - 1p Lunch	12p - 1p Lunch	11:15 - 12p Dealing with Dynamic Data		
12pm	1p - 2:30p Loop Blocking and Register Tiling for Locality	1p - 2:30p Keynote 1: Algorithm Design for GPU Computing	1p - 2:30p Dealing with Non-uniform Data Distribution	12p - 1p Lunch		
1pm	2:30p - 3:30p Lab 1	2:30p - 3:30p Lab 2	2:30p - 3:30p Lab 3	2p - 2:30p Lab 3		
2pm						
3pm						
4pm						

Events shown in time zone: Pacific Time

Google Calendar

Course Presentations - VSCSE Many-core Processors 2010 | Google Groups

http://groups.google.com/group/vscse-many-core-processors-2010/web/course-presentations

Google Groups will no longer be supporting the Pages and Files features. Starting January 13, you won't be able to upload new content, but you will still be able to view and download existing content. See [this announcement](#) for more information and other options for storing your content.

Course Presentations

See the [course schedule](#) for the precise times of each presentation. Speaker [biographies](#) are posted on a separate page.

Monday, August 2

- Lecture 1: Introduction to Computational Thinking for Many-core Computing (Wen-mei Hwu, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 2: Scatter-to-Gather Transformation for Scalability (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 3: Loop Blocking and Register Tiling for Locality (Wen-mei Hwu, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Tuesday, August 3

- Lecture 4: Cut-off and Binning for Regular Data Sets (Wen-mei Hwu and John Stone, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 5: Data Layout for Grid Applications (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Keynote 1: Algorithm Design for GPU Computing (Michael Garland, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Wednesday, August 4

- Keynote 2: Fermi and the Future of GPU Computing Technology (David Kirk, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 6: PDE Solver Techniques (Jonathan Cohen, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 7: Dealing with Non-uniform Data Distribution (Wen-mei Hwu, University of Illinois; David Kirk, NVIDIA) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Thursday, August 5

- Keynote 3: Multiplying speeds: fast algorithms on GPUs (Lorena Barba, Boston University) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Guest Lecture: Accelerating HPC Applications with GPUs—Two Case Studies (Jeremy Meredith, Oak Ridge National Laboratory) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Lecture 8: Dealing with Dynamic Data Sets (Wen-mei Hwu, University of Illinois; David Kirk, NVIDIA) [\[video-part-1\]](#) [\[video-part-2\]](#) [\[ppt\]](#) [\[pdf\]](#)
- Hands-on Lab Wrap-up Discussion (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Friday, August 6

- Lecture 9: Directions for Further Studies and Closing Remarks (Wen-mei Hwu and John Stratton, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)
- [gpubcomputing.net](#) (Andrew Schuh, University of Illinois) [\[video\]](#) [\[ppt\]](#) [\[pdf\]](#)

Home

Discussions

Pages

Files

[About this group](#)

[Apply for group membership](#)

Sponsored links

GPU Computing Cluster
HPC Cluster with Intel® Xeon® CPUs and Tesla GPUs. Request a Quote!
[www.Appro.com](#)

University of Phoenix®
Online and Campus Degree Programs.
Official Site - Classes Start Soon.
[Phoenix.edu](#)

Group Info

Members: 158

Language: English

Group categories: Not categorized

[More group info >](#)

Recent pages and files

<input type="checkbox"/> Course Presentations	Aug 24
<input type="checkbox"/> Participating Sites	Aug 10
<input type="checkbox"/> Hands-on Labs	Aug 9
<input type="checkbox"/> Teaching Assistant	Aug 9
<input type="checkbox"/> Solutions	Aug 6

Participating sites

Gmail Calendar Documents Reader Web more ▾

Google groups

VSCSE Many-core Processors 2010

Search this group Search Groups

Google Groups will no longer be supporting the Pages and Files features. Starting January 13, you won't be able to upload new content, but you will still be able to view and download existing content. See [this announcement](#) for more information and other options for storing your content.

Participating Sites

The following sites are fully participating in the *Proven Algorithmic Techniques for Many-core Processors* course:

- Center for Computation & Technology, Louisiana State University, Baton Rouge [\[local info\]](#) [\[photos\]](#)
- Institute for Data and High Performance Computing, Georgia Institute of Technology, Atlanta [\[local info\]](#) [\[photos\]](#)
- Institute for Digital Research and Education, University of California, Los Angeles [\[local info\]](#) [\[photos\]](#)
- National Center for Supercomputing Applications, Urbana, Illinois [\[local info\]](#) [\[photos\]](#)
- Northwestern University, Evanston, Illinois [\[local info\]](#) [\[photos\]](#)
- Ohio Supercomputer Center, Ohio State University, Columbus [\[local info\]](#) [\[photos\]](#)
- RENCI, Chapel Hill, North Carolina [\[local info\]](#) [\[photos\]](#)
- University of Iowa, Iowa City [\[local info\]](#) [\[photos\]](#)
- University of Michigan, Ann Arbor [\[local info\]](#) [\[photos\]](#)

The following sites will host remote presenters (with no audience):

- University of Minnesota Supercomputing Institute (Minneapolis)
- NVIDIA Corporation (Santa Clara, CA)

The following site is a receiving site only:

- Lawrence Berkeley National Laboratory [\[local info\]](#)

These sites are linked by a [high-definition videoconferencing technology](#), which provides a collaborative learning environment spanning geographical boundaries.

Home

Discussions

Pages

Files

[About this group](#)

[Apply for group membership](#)

Sponsored links

GPU Computing Cluster
HPC Cluster with Intel® Xeon® CPUs and Tesla GPUs. Request a Quote!
[www.Appro.com](#)

IVR made *dam* easy
Voxeo XML IVR platforms & hosting:
Free trials, extremely good support
[www.voxeo.com](#)

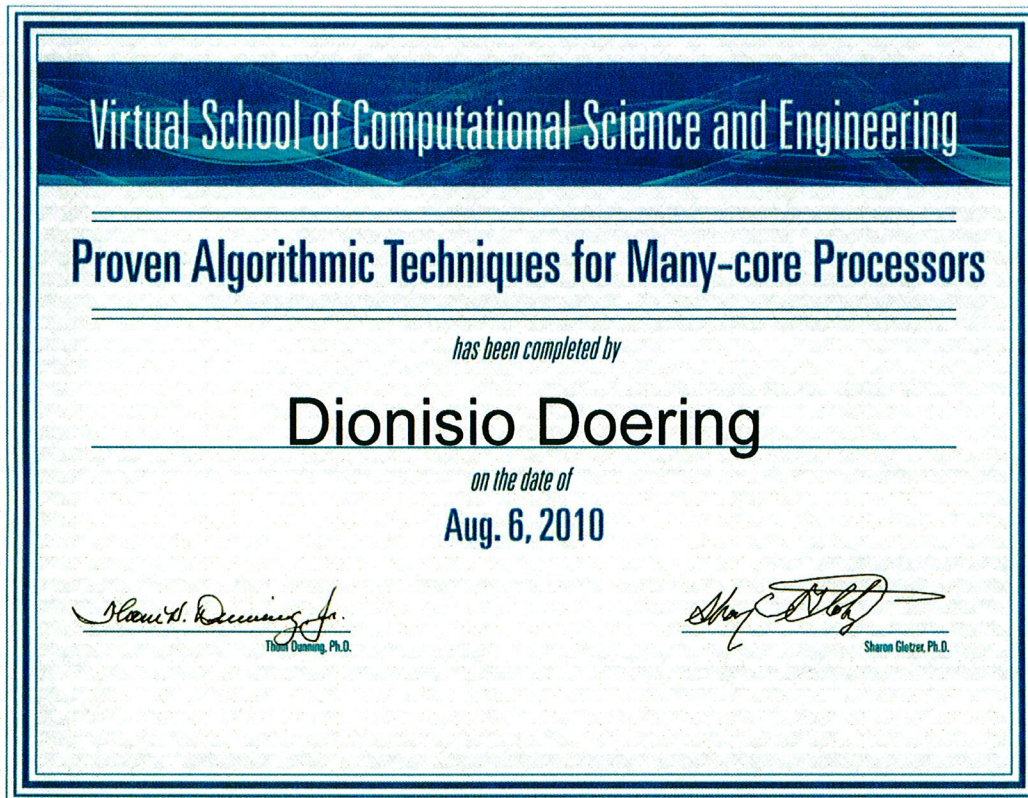
San Francisco Coupons
1 ridiculously huge coupon a day.
Get 50-90% off the Bay's best
[www.Groupon.com/San-Francisco](#)

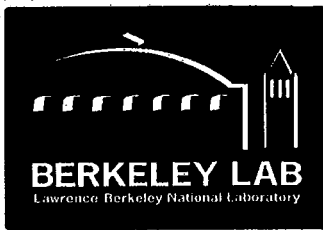
Group Info

Members: 158

Version: Current by Tom Scavo - Aug 10

Course certificate





ENGINEERING DIVISION
Advancing Science By Design

1 Cyclotron Road, MS: 46A1123
Berkeley, CA 94720
Phone: 510.486.4200 Fax: 510.496.2403

December 9, 2010

Programa de Pós-Graduação em Engenharia Elétrica - PPGEE
Escola de Engenharia
Universidade Federal do Rio Grande do Sul - UFRGS
Av. Osvaldo Aranha, 103
90035-190 Porto Alegre, RS - Brasil

To whom it may concern:

This letter is to confirm that we are aware that Dionisio Doering is enrolling into your Ph.D. program as a part time student. The Engineering Division at Lawrence Berkeley National Laboratory encourages the continuation of education for all Laboratory employees and we are aware that his enrollment in this program will not impact his work hours.

Sincerely,

Dr. Peter Denes
Acting Engineering Division Director

cc: Dionisio Doering



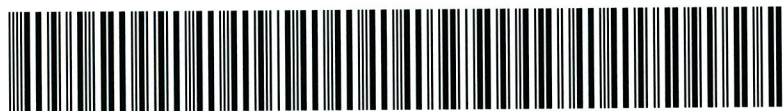
MINISTÉRIO DA FAZENDA
SECRETARIA DO TESOURO NACIONAL
Guia de Recolhimento da União

RECIBO DO SACADO

CEDENTE UNIVERSIDADE FEDERAL DO RGS - CNPJ: 92.969.856/0001-98				NOSSO NÚMERO 10812640000198119	VENCIMENTO 17/12/2010
DATA 10/12/2010	NÚMERO DOC 198119	AGÊNCIA/CEDENTE 3798-2 / 333002-8	(=) VALOR R\$ 100,00	(+) MORA E MULTA R\$ 0,00	(=) VALOR COBRADO R\$ 100,00
OBSERVAÇÕES: Contribuinte: Dionisio Doering CGC/CPF: 916.090.890-49 Recolhimento: Taxa de Inscrição Concurso e Processos Seletivos Orgão: Escola de Engenharia Observações Gerais: Inscrição Doutorado 20111 Ppgee				AUTENTICAÇÃO MECÂNICA	

BANCO DO BRASIL | **001-9** | 00190.00009 01081.264002 00198.119182 7 48190000010000

LOCAL DO PAGAMENTO ATÉ O VENCIMENTO, PAGÁVEL EM QUALQUER AGÊNCIA DA REDE BANCÁRIA					VENCIMENTO 17/12/2010
CEDENTE UNIVERSIDADE FEDERAL DO RGS - CNPJ: 92.969.856/0001-98					AGÊNCIA / CÓDIGO CEDENTE 3798-2 / 333002-8
DATA DO DOCUMENTO 10/12/2010	NÚMERO DO DOCUMENTO 198119	ESPÉCIE DOC RC	ACEITE N	DATA PROCESSAMENTO 10/12/2010	NOSSO NÚMERO 10812640000198119
USO DO BANCO 153301	CARTEIRA 18/19-1	ESPÉCIE R\$	QUANT. 1	VALOR	(=) VALOR DO DOCUMENTO R\$ 100,00
INSTRUÇÕES: NÃO RECEBER APÓS A DATA DE VENCIMENTO. Recolhimento: 28883-7 - Taxa de Inscrição Concurso e Processos Seletivos Orgão/UGR: 153301 - Escola de Engenharia Observações Gerais: Inscrição Doutorado 20111 Ppgee					(-) DESCONTO
					(-) OUTRAS DEDUÇÕES
					(+) MORA E MULTA
					(-) OUTROS ACRÉSCIMOS
					(=) VALOR COBRADO R\$ 100,00
SACADO/CONTRIBUINTE Dionisio Doering 916.090.890-49					



AUTENTICAÇÃO MECÂNICA

FICHA DE COMPENSAÇÃO

Pagamento de boletos bancários e títulos na fatura do cartão

10/12/2010 - BANCO DO BRASIL - 15:56:52
189901899 0007
OUVIDORIA BB 0800 729 5678

AGENDAMENTO DE PAGAMENTO DE TITULOS

CLIENTE: DIONISIO DOERING
CARTAO.: 4984 **** * 4239

=====

BANCO DO BRASIL

00190000090108126400200198119182748190000010000
NR. DOCUMENTO 121.701
NOSSO NUMERO 10812640000198119
CONVENIO 01081264
UNIVERSIDADE FEDERAL DO RIO GR
AGENCIA/COD. CEDENTE 3798/00333002
DATA DE VENCIMENTO 17/12/2010
DATA DO PAGAMENTO 17/12/2010
VALOR DO DOCUMENTO 100,00
VALOR COBRADO 100,00

=====

PAGAMENTO AGENDADO.
PAGAMENTO COM CARTAO DE CREDITO A VISTA
SUJEITO A COBRANCA DE TARIFA DE ACORDO COM A
TABELA DE TARIFAS VIGENTE NA DATA DA LIQUIDACAO.

Transação efetivada com sucesso!

Evite a impressão dos seus comprovantes utilizando a opção
SALVAR COMPROVANTE. O meio ambiente agradece.

